

12 **EUROPEAN PATENT APPLICATION**

21 Application number: 85401029.5

61 Int. Cl.⁴: G 06 F 15/16

22 Date of filing: 24.05.85

30 Priority: 01.06.84 US 616553

43 Date of publication of application:
04.12.85 Bulletin 85/49

84 Designated Contracting States:
BE CH DE FR GB IT LI NL SE

71 Applicant: **DIGITAL EQUIPMENT CORPORATION**
148 Main Street
Maynard, MA 01754(US)

72 Inventor: **Mann, Bruce**
Valley Road
Mason New Hampshire 03048(US)

72 Inventor: **Duffy, Darrell**
3 Nashua Road
Windham New Hampshire 03087(US)

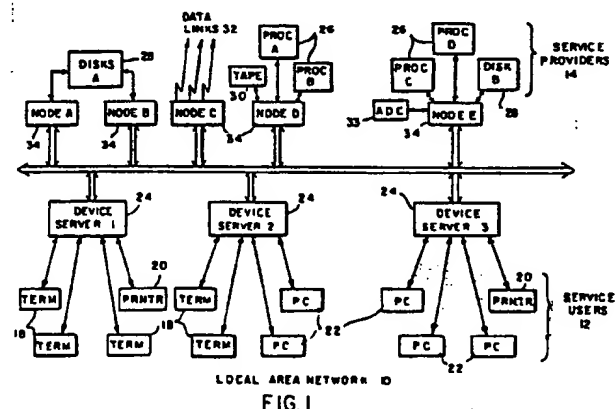
72 Inventor: **Lauck, Anthony**
20 Fells Circle
Wellesley Massachusetts 02187(US)

72 Inventor: **Strecker, William**
33 Ann Lee Road
Harvard Massachusetts 01451(US)

74 Representative: **Mongrédién, André et al,**
c/o SOCIÉTÉ DE PROTECTION DES INVENTIONS 25, rue
de Ponthieu
F-75008 Paris(FR)

64 Local area network for digital data processing system.

57 A local area network for interconnecting terminals and other users and data processing systems and other service providers over a communications link. The users and providers connect to the communications link by means of interface units each of which may connect to several users or providers. The interface units communicate over the communications link by means of messages. When a user requires the use of a service, the interface unit establishes a virtual circuit between it and the interface unit connected to the service provider and a service session which allows the user and the service provider to communicate over the virtual circuit. If several users connected to the one interface unit as the first user require services provided by providers which connected to the same interface unit as the first provider, they communicate in sessions over the same virtual circuits. The sessions messages are accumulated into single virtual circuit messages that are acknowledged in unison by the receiving interface unit. Each virtual circuit in the users' interface units includes a timer which reset when a message is transmitted over the virtual circuit and a data waiting flag set whenever data is present to be transmitted over the virtual circuit. The interface units are inhibited from transmitting over a virtual circuit unless the timer has timed out and the data waiting flag is set.



Background of the Invention

1. Field of the Invention

The invention relates generally to the field of digital data processing systems and more specifically to local area networks in which a plurality of such systems are interconnected to provide distributed processing capabilities to a number of users. In particular the invention provides improvements in message transfer protocols for local area networks to enhance the message transfer capability of the network.

2. Description of the Prior Art

As small and medium-sized computer systems are becoming less expensive and more powerful, a number of them are being interconnected to form networks to ensure that a number of different types of services are available at any time to users having diverse processing needs. Such services may include any of the services which are normally available from such networks, including such as electronic mail (mail storage and forwarding), word processing, accounting, such as payroll or inventory, or data communications over telephone lines or microwave links. Interconnecting systems into a network helps to enhance the availability of services to service users by including a number of service providers in the network and having each provider provide one or more of the services, thus reducing the likelihood that the failure of any one service provider in the network will result in a significant number of services being unavailable to the users at any one time. Indeed, a local network may be arranged so as to have several service providers providing overlapping services, in which case several providers have the ability to provide

a particular service if one service provider fails or is saturated with service requests.

Typically in a local area network, the communications in the network take place over one or a limited number of communications links. Examples of such communications links include those defined by well-known DECnet, SNA (System Network Architecture) or X.25 communications protocols using data links such as Ethernet. A number of service users, such as, for example, video terminals controlled by operators, are connected through interface devices known as "terminal servers" to the communications link. Similarly, the service providers are connected to the communications link through interface devices known as "nodes".

If the operator desires to use a service provided by a unit connected to a node, it may request connection to the node and, through the node, to the unit to have the service provided. Normally, the operator has to know the particular node(s) and unit(s) that provides the desired service. The operator selects a node and unit to provide the service, and causes the terminal server to request service by that node and service provider. The terminal server and the node exchange messages which enable a "virtual circuit" to be established which provides a data transfer mechanism between the operator's terminal and the provider of the service. The virtual circuit essentially extends from the operator's terminal, as the service user, through the terminal server, over the communications link and through the node to the service provider. If a number of users are using the local area network, several virtual circuits may be established over the communications link to provide communications between the users and providers. In addition, if several terminals connected to one terminal server require

services from a service provider connected to the same node, separate virtual circuits are normally established between each terminal and service provider providing the required service.

5 Service data is transmitted in the form of messages through the virtual circuits between the terminals and the service providers. All of the messages are queued by the terminal server and transmitted over the single communications link. To ensure that the messages are
10 received, units connected in the network re-transmit the message until an acknowledgement from the recipient is received verifying correct receipt. Specifically, the terminal servers and nodes, after they transmit messages through the virtual circuits over the communications
15 links, monitor the communications link for acknowledgements and if no acknowledgement is received within a selected period of time, corrective action is taken. Each message transmitted through a virtual circuit is acknowledged by a separate acknowledgement
20 message through the virtual circuit, even if a series of messages are transmitted between the same terminal server and node through different virtual circuits over the same communications link. Each acknowledgement message must be separately generated, and thus requires that time and
25 facilities be dedicated at the receiving device to the generation of the separate acknowledgement messages. Furthermore, requiring such separate acknowledgement messages could cause the communications link to quickly become unnecessarily burdened.

30 In current networks, message transfers through virtual circuits over a communications link are initiated either by the occurrence of certain events ("event driven" transmission), such as the presence of data to be transmitted, or by the timing-out of certain timers

("timer-based" transmission). Both the event-driven and timer based message transfer systems incorporate certain assumptions about message traffic through virtual circuits over the communication link. The event driven systems assume that the communications link has sufficient bandwidth, even when it is being heavily used, to ensure that messages can be delivered from the terminal servers to the nodes, the messages can be processed by the service providers, and the responses to the messages can be returned to the terminal servers, all within a maximum tolerable delay period. If the communications link carries too high a level of message traffic, the delays will, however, become unacceptably long. Furthermore, if a unit transmits data every time it receives a few bytes of data from a service user or provider, a significant number of such messages will be composed primarily of virtual circuit identification information, which is necessary to ensure proper identification of the virtual circuit carrying message over the multiplexed communications link but otherwise serves no purpose.

Timer-based message transmission systems ensure, on the other hand, that every unit connected to the communications link will be able to transmit messages over the link at periodic intervals. These systems ensure that all of the units have relatively uniform access to the communications link. However, such systems also have a number of deficiencies. First, each unit, when its time comes to transmit, transmits messages through its virtual circuits whether or not it has any data to transmit, obviously wasting bandwidth on the communications link. Furthermore, as units are added to the system, the timers of all of the units would have to

be adjusted to ensure that all of the units have reasonably equal access to the communications link.

Summary of the Invention

5 The invention provides a new local area network message transfer system which has enhanced message throughput between service users and service providers over a communications link in a local area network, while at the same time ensuring that all units have a
10 relatively uniform access to the communications link for transferring messages.

 In brief summary, the invention provides a local area network over which a plurality of users, such as terminals or printers, communicate with service
15 providers, such as data processing systems, data storage devices such as disk or tape records, or data links such as telephone lines or microwave transmission links. One or more device servers connects directly to a communications link. Each device server is an interface
20 to the communications link for one or more of the users. Similarly, one or more nodes connects to the communications link, and each node is an interface to the communications link for one or more service providers. Periodically, each node transmits a service advertising
25 message over the communication link which includes its identification, the identifications of the services that are provided by the service providers connected to it, and a ratings for each of the services. Each device server receives these messages and records them in a
30 service directory. The services available to a user by the service providers in the network may be viewed by the operators from the service directory stored by the device servers.

When an operator desires to use a service in the service directory, it enables the device server to request the service. The device server then selects the particular provider to provide the service based on the ratings in the service directory, and identifies the node through which communications with the service provider can be conducted. If the device server is not then communicating with that node, it and the node establish a virtual circuit through which they transfer messages. In addition, the device server establishes a service session between itself the user whose operator is requesting the service, and the node establishes a service session between itself and the service provider connected to that node which provides the service requested by the user, with the service sessions at the device server and node, respectively, being linked and the identification of the session being known to the unit, that is, the device server or node, at the other end of the virtual circuit. If any other users request services provided by service providers connected to that same node, similar sessions are established for those users by the device server and for the service provider by the nodes, and messages for all such sessions may be transferred through the same virtual circuit multiplexed in slots in the same virtual circuit message which serves all of the service users or providers. Therefore, the device server need not establish a new virtual circuit for every device which requires the service provided by a service provider connected to that node. Furthermore, only the virtual circuit messages are acknowledged, rather than the individual messages between service users and providers, thereby reducing the number of acknowledgement messages transferred and the resources at the device servers and nodes required to generate the acknowledgement messages.

Generally, message transfers between a device server and a node through a virtual circuit are initiated by a device server, and each message from a device server to a node is acknowledged by a message from the node. Each
5 message includes a response requested flag which may be set or cleared by the node. The response requested flag is set if the message includes session slot data, otherwise the flag is cleared. The portion of each virtual circuit in the device server includes a server
10 circuit timer and a data waiting flag (DWF) which is set either by the receipt of a message from the node having a set response requested flag or in response to the receipt of slot data from the service users which use the virtual circuit. Normally, the node will send a message only in
15 response to a message from the device server; however, if the response requested flag in the prior message was cleared, which occurs if the node did not send any data in the prior message, the node may send another message to the device server which includes data before it
20 receives a subsequent message from the device server. When a device server transmits a message through the virtual circuit over the communications link, the server circuit timer, pre-set to a selected value, begins to decrement. When the server circuit timer times out and
25 if the data waiting flag is set, the device server transmits a message through the virtual circuit. Until the server circuit timer times out, the device server is inhibited from transmitting a new message over the virtual circuit. Furthermore, the data waiting flag
30 inhibits the device server from transmitting any messages until the data waiting flag is set indicating it has received a message from the node in which the response requested flag is set or that it has new slot data from the service users to send.

The response requested flag allows the node to immediately acknowledge a message received from the device server, whether or not it has data to send immediately, and to reserve for itself the ability to transmit a second message which has data if no data is sent in the initial message, which second message can be sent even if it receives no intervening message from the device server. The response requested flag in conjunction with the data waiting flag allows the node to force the device server to acknowledge a message that includes data when the server circuit timer next times out, whether or not the device server has data to send. The server circuit timer establishes a minimum delay period after the device server transmits a message before it transmits a subsequent message, thereby allowing other device servers to transmit messages through virtual circuits which they have established over the communications link. The device server and the node thus have the benefit of both a timer based system, based on the server circuit timer, and an event driven system, based on the response requested and data waiting flags.

Brief Description of the Drawings

This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following detailed description taken in conjunction with the accompanying drawings in which:

FIG. 1 is a general block diagram of a local area network constructed in accordance with this invention;

FIG. 2A is a diagram illustrating the contents of a service advertising message transmitted by the service providers in the network depicted in FIG. 1, and FIG. 2B is a diagram of a data base established by the service

users in the network of FIG. 1 in response to the service advertising message depicted in FIG. 2A;

FIG. 3A is a diagram depicting a virtual circuit and service sessions that is useful in understanding the operations of the network depicted in FIG. 1;

FIGS. 3B and 3C are diagrams depicting databases used by the service providers and service users in the network depicted in FIG. 1;

FIG. 4 is a state diagram useful in understanding the operation of the virtual circuit depicted in FIG. 3A;

FIG. 5 is a state diagram useful in understanding the operation of the service sessions depicted in FIG. 1;

FIGS. 6A through 6D depict the formats of virtual circuit messages transmitted through the virtual circuit depicted in FIG. 3A;

FIGS. 7A through 7D depict the formats of session slot messages transmitted between corresponding sessions in the device servers and nodes as depicted in FIG. 3A; and

FIG. 8 depict the timings of messages transmitted through the virtual circuits as depicted in FIG. 3A.

Detailed Description of an Illustrative Embodiment

FIG. 1 illustrates a local area network 10 in which a plurality of service users, generally indicated by reference numeral 12, communicate with a plurality of service providers, generally indicated by reference numeral 14, over a common communications link 16. A communications link 16 may take the form of any one of a number of communications lines and interface circuitry which transfer data between the service users and service providers in bit serial or parallel form. For example, the communications link may take the form of a coaxial cable and interface circuitry which transmits messages using the well-known Ethernet local area network

protocol. In that protocol, data is organized into messages having a predetermined format and transmitted in bit serial form between stations over the coaxial cable. A number of other communications links using diverse
5 protocols exist which could also be used in the local area network depicted in FIG. 1; the specific communications link selected is not an aspect of the invention.

The service users 12 may include a plurality of
10 devices such as, for example, video display terminals 18, printers 20, and personal computers 22. Network 10 also includes a plurality of device servers 24 each of which connect to several service users and enable the service users to communicate over communications link 16 with the
15 service providers 14. The service providers 14 may include devices such as processors 26, disk drives 28, tape storage units 30 and data links 32 (such as telephone lines or microwave links) and analog-to-digital converters 33. Network 10 also includes a plurality of
20 nodes 34, each of which may be connected to several service providers. Furthermore, each node 34 connects directly to communications link 16 and provides communications between the service providers connected to it and the communications link. It will be appreciated
25 that in some cases a service provider and node may be integrated into one unit that performs both functions.

The service providers 14 provides service to the service users 12. Such services may include, for example, electronic mail storage and forwarding among
30 service users, word processing capabilities, access to programs such as payroll accounting, inventory control or the like, the ability to store or retrieve records on or from disk and tape files, the ability to communicate over telephone lines and microwave links, and the ability to

acquire data from, for example, scientific instruments through analog-to-digital converters 33. Such services, as well as additional services, are well known in the art and will not be discussed further herein.

5 It will be appreciated by those skilled in the art that some service users may also provide services. For example, certain personal computers 22 in network 10, in addition to being a service user, may also have programs that may be accessed and used by another user, such as a
10 terminal 18. In that case, the personal computer may be connected to a node 34 as well as to a device server 24 to make its programs available to a service user as services. A unit which interfaces the personal computer to the communications link may provide the facilities of
15 both a node and a device server.

Periodically, each node 34 transmits an "advertising" message which is received by all of the device servers 24, which identifies itself and the services provided by the service providers 14 connected
20 to that node, as well as a "rating" for each service. With reference to FIG. 2A, the service advertising message includes a plurality of fields including a header 50 which depends on the protocol used over a communications line 16 and a body 54. In one embodiment,
25 the header 50 has a node identification field 51 which identifies the transmitting node, a protocol type field 52 identifying the message as the service advertising message, and a multi-cast address field 53
30 which enables all of the device servers 24 to receive the message. Following the header 50, the node transmits the body 54 of the message which identifies the various services provided by it and the rating for each service. The ratings indicate, for example, how prompt the service provider 14 may be in responding to a service request,

based in part on the number of service users 12 then using the service provided by the specific provider and, thus, the potential delay in responding to communications from another service user who might request the service.

5 After receiving service advertising messages from the nodes 34, each device server 24 (FIG. 1) establishes a service directory such as depicted in FIG. 2B. The directory comprises a table in which the node identifications, the services provided by the nodes, and
10 the service ratings are stored. Thus, if an operator at a service user 12 desires to use one of the services shown in the directory, the device server can use the contents of the service directory depicted in FIG. 2B to determine which nodes provide that service. If more than
15 one node provides the service requested, the device server uses the rating in the rating fields to determine which node has the highest rating for the service and requests that node to provide the service.

 The diverse services that are available from the
20 service providers may be divided into groups or classes and each user may be able to access only those services that are of interest to him. The service names may be organized into groups identified by group names, and the device server will display only those services to a
25 particular user which are in the groups which the user can access.

 When a user 12 requires a service provided by a service provider 14 identified in the service directory, the device server 24 begins to establish a virtual
30 circuit 58 between itself and the node 34 that provides the service with the most desirable service rating. With reference to FIG. 3A, the device server 24 in a conventional manner establishes a virtual circuit state machine 60 which provides two-way data communications

over a pair of unidirectional data pipes with a virtual circuit state machine 64 established by node 34. The virtual circuit state machines 60 and 64 and the data pipes 62 provide a means for transferring data, in the form of messages between the device server 24 and the node 34 over the communications link 16. It will be appreciated that message communications through a number of data pipes 62 may be multiplexed over the communications link 16, and, accordingly, the communications link provides message communications for a number of virtual circuits in network 10.

The virtual circuit state machine 60 at the device server 24 communicates with the individual service users 12 by means of service sessions using separate session state machines 66 which the device server establishes in a conventional manner for each user. Similarly, the node's virtual state machine 64 communicates with the service providers 14 using separate session state machines 68.

The device server 24 and node 34 use messages transmitted over communication link 16 to set up the virtual circuit and the session state machines, which will be described below in connection with FIGS. 4 through 7D. In brief, however, when a user 12 requires service by a service provider 14, the device server 24 first determines whether a virtual circuit exists between it and the node 34 selected by the device server. If no such virtual circuit exists, the device server 24 transmits a virtual circuit message over communications link 16 to node 34 which causes the node to establish its virtual circuit state machine 64 to support its end of the virtual circuit 58. A session state machine 66 is also set up between its virtual circuit state machine 60

and the user 12 to allow data and other information to be accumulated from and transferred to the user.

In a succeeding virtual circuit messages, after the virtual circuit is set up, a session slot is transmitted
5 by device server 24 to node 34 through the virtual circuit 58, specifically over the communications link 16, identifying the required service, and the node 34 sets up a session state machine 68, which connects to the service provider which provides the required service and allows
10 data and other information to be transferred between the virtual circuit state machine and the service provider. Each session state machine collects information to be transferred from its connected user or service provider and provides the information in the form of session
15 messages to the virtual circuit state machine, which in turn accumulates the session messages from various service users' or providers' state machines which are to be transferred between the same device server 24 and node 34 and forms a single virtual circuit message for
20 transfer through the virtual circuit 58 over the communications link 16. On receiving a virtual circuit message from the virtual circuit 58, the receiving virtual circuit state machine transfers the session messages to the respective session state machines that
25 are the intended recipients for transfer to the respective service users 12 and service providers 14, and returns a single acknowledgement message over the virtual circuit to verify receipt of the virtual circuit message. It will be appreciated that requiring only a single
30 virtual circuit acknowledgement message for the multiplexed messages between service users and providers reduces the acknowledgement message traffic from that often required in the prior art, thereby reducing the traffic overhead over communications link 16, and also

reduces the overhead required at the device server and node to generate the acknowledgement messages.

The virtual state machines 60 and 64 in device server 24 and node 34 respectively include a data base as depicted in FIG. 3B used in transmitting and receiving messages through pipes 62. Data base 70 includes a remote identification word 72 and a local identification word 74. The identification words 72 and 74 contain the identification of the virtual circuit 58 as assigned by node 34 and device server 24. The contents of the local identification word 74 are assigned by the unit in which the data base 70 resides, and the contents of the remote identification word 72 are assigned by the other unit engaged in the virtual circuit. Thus, in the virtual circuit data base 70 which resides in device server 24, the local identification word is assigned by the device server and the remote identification word is assigned by the node 34 which provides the other end of the virtual circuit. Similarly, in the virtual circuit data base 70 (FIG. 3B) residing in node 34, the contents of the local identification word 74 are assigned by the node, and the contents of the remote identification word 72 is assigned by the device server. The contents of the two identification words 72 and 74 are transmitted in the virtual circuit messages transferred through the virtual circuit over communications link to allow the device server and node to identify the messages transferred over communications link 16 as being associated with the particular virtual circuit.

The virtual circuit data base 70 also includes a message type field 76 which identifies the type of virtual circuit message to be transmitted next. Three types of virtual circuit messages are transmitted over a virtual circuit, namely, a START virtual circuit message,

a RUN message, and a STOP virtual circuit message, which are detailed below in connection with FIGS. 4 and 6A through 6D.

An M field 78 identifies whether the unit including
5 the data base 70 is a master or slave unit. In the network 10 (FIG. 1) device servers 24 are always masters and nodes 34 are always slaves; communications over a virtual circuit are always initiated by the device servers and the nodes always acknowledges or responds to
10 the communications from the device servers.

An R field 80, when set, indicates that the last sent message requires a response.

Data base 70 also includes message counters 82 and acknowledgement counters 84. Each message transmitted by
15 device server 24 or node 34 includes a message sequence number that is checked by the virtual circuit state machine in the receiving unit to ensure that the successive messages are properly received in sequence. Messages are re-transmitted to ensure that they are
20 properly received if they are not acknowledged within a predetermined time; the sequence numbers ensure that the receiving device does not treat a re-transmitted message as a new message if it, in fact, correctly received the message on a previous transmission. The contents of the
25 message counters 82 identify the number of messages which have been transmitted and received, and the acknowledgement counters contain the message numbers of the messages which have been acknowledged. Thus, if a message number is skipped, or if acknowledgements are not
30 received in numerical order, the device server and node can determine which messages may have not been properly transmitted over communications link 16.

A data waiting flag (DWF) 86 is set whenever any session state machine has data to send over the virtual

circuit. In a device server 24, the data waiting flag is set when a message is received from a node which requires a response.

5 A retransmit counter 88 and retransmit timer 90 are used in retransmitting messages which have not been acknowledged within a time selected by the retransmit timer. The transmitting unit retransmits each unacknowledged message a number of times as selected by the retransmit counter 88. If the message is not
10 acknowledged after the retransmit counter counts out, the other end of the virtual circuit is marked out of service.

The virtual circuit data base 70 in a device server 24 also includes a server circuit timer 92. When
15 the device server 24 transmits a message, it resets its server circuit timer 92 and is thereafter inhibited from transmitting a subsequent message until the server circuit timer times out. Thus, even if the data waiting flag 86 is set, indicating that the device server has
20 information to transmit in a virtual circuit message, the device server is inhibited from transmitting virtual circuit messages over virtual circuit 58 until the server circuit timer times out. Conversely, even if the server circuit timer 92 has timed out the device server 24 does
25 not transmit any messages unless the data waiting flag 86 is set. After the server circuit timer times out, if the virtual circuit's data waiting flag 86 is thereafter set and the previously transmitted messages have been .
acknowledged, the device server 24 may then immediately
30 transmit a new virtual circuit message over the virtual circuit 58. If a previously transmitted message has not been acknowledged, a device server or node waits an amount of time as specified by the retransmit timer 90 and then retransmits the unacknowledged message.

The server circuit timer 92 thus ensures at least a minimum delay period between the transmission of new virtual circuit messages by the device server 24 through any particular virtual circuit. Thus, when traffic is heavy through the virtual circuit and the data waiting flag is set before the server circuit timer times out, message transmissions will be based on the timing out of the server circuit timer 92. However, if traffic is light through the virtual circuit, the virtual circuit message transmission will be based on the setting of the data waiting flag 86. Since, with one exception as explained below, the nodes 34 only respond to or otherwise acknowledge virtual circuit messages from the device servers, the server circuit timers and the data waiting flags in the device servers also govern message traffic from the nodes through the virtual circuits. The network 10 (FIG. 1) thus achieves the benefits of both the timer-based message transfers and the event-driven message transfers, the event-driven transfers being initiated by the presence of information from user 12 for transfer over the virtual circuits, and the timer-based message transfers being based on the timing out of the server circuit timer, without having the detriments of either.

The virtual circuit data base 70 in a device server 24 also includes a keep-alive timer 94 which enables the device server to send a message over the virtual circuit if it has not sent any messages thereover for a very long period of time, to ensure that the node 34 maintains support for its end of the virtual circuit. The node 34 responds thereby informing the device server that it has not gone down.

As noted above, the service user and service provider providing the service required by the user

communicate by means of session slots. More specifically, the session state machines at the device server and node transfer session slots which cause transitions between states in the respective session state machines and also transfer service data and status information between the service user and provider.

Each session state machine uses a session data base 100 depicted in FIG. 3C. The session data base includes a remote identification field 102 and a local identification field 104, which are used in the same way as a virtual circuit state machine uses the remote and local identification fields 72 and 74 in the virtual circuit data base 70 (FIG. 3B). Specifically, each virtual circuit message that is transmitted over a virtual circuit may include session slots for different service sessions (that is, session slots intended to be used by different session state machines in the device servers and nodes that communicate over the same virtual circuit) and the remote and local identification fields 102 and 104 identify the session and session state machines that are the intended recipients of the session slots. The contents of local identification field are assigned by the unit in which the data base resides, and the contents of the remote data base are assigned by the other unit.

Each session data base 100 also includes a data buffer 106 for storing data received from or to be transmitted to the service user 12 or service providers 14 associated with the particular session state machine.

When the data buffer 106 is loaded with data from the user 12 or service provider 14 connected to the particular session state machine, a data ready flag 108 is set, which in turn enables the data waiting flag 86 in

the virtual circuit data base 70 to be set. When the device server or node thereafter sends a message virtual circuit, it can poll the data ready flags of the service sessions assigned to the virtual circuit to determine
5 whether their data buffers have data to transmit, and may remove the contents of various fields, including the data from the data buffers 106 whose associated data ready flags as set, to generate session slots for transfer in a virtual circuit message.

10 A byte count field 110 identifies the number of bytes of data in the data buffer 106 and is sent with the data in the session slot.

A session slot type field 112 identifies the type of session slot to be sent. Five types of session messages
15 can be sent, including START, STOP, REJECT, DATA, and STATUS messages. The contents of the session messages will be described below in connection with FIGS. 5 and 7A through 7D.

Local and remote credits fields 114 and 116 in the
20 session data base 100 relate to the number of slots that are available, with each slot relating to a specific amount of data. Each session slot transmitted over a virtual circuit includes a credits field which identifies the amount of space available in the data buffer for any
25 response information from the unit at the other end of the virtual circuit engaged in the service session. The contents of the credits field in the message is provided by the contents of the local credits field 114 in the session data base 100 of the unit transmitting the
30 message. The contents of the remote credits field 116 is provided by the contents of the credits portion of the session slot received from the unit at the other end of the virtual circuit.

FIG. 4 illustrates the various states of virtual circuit state machines 60 and 64 (FIG. 3A) and the various messages which can be transmitted over the virtual circuit during those states and which cause transitions among states. FIGS. 6A through 6D detail the contents of the various virtual circuit messages. As noted above, three types of virtual circuit messages are transmitted through the virtual circuit, including a START virtual circuit message, a STOP virtual circuit message, and a RUN message. The START and STOP virtual circuit messages are used to establish and abolish the virtual circuit, and the RUN message is used to transfer information, including session slots, between service users and providers.

With reference to FIG. 4, the state machines 60 and 64 in the device server 24 and node 34 respectively are both initially in a HALTED state. When a user 12 requires a service provided by a service provider 14 connected to a node 34, if no virtual circuit exists between the device server 24 and that node 34, the device server 24 transmits a START virtual circuit message to node 34.

FIG. 6A depicts the general format of a virtual circuit message. With reference to FIG. 6A, the message begins with a communications link header 120, the format of which depends on the particular communications link 16 selected for the network 10. In the specific embodiment in which the communications link conforms to the Ethernet protocol, the communications link header includes a destination address field 122 and a source address field 124 which identify the particular sending and receiving node and device server, and a protocol type 126.

After the communications link header 120, the message includes a virtual circuit header 130 which identifies the virtual circuit over which the message is being transmitted. The virtual circuit header includes a
5 destination virtual circuit identification field 132 and a source virtual circuit identification field 134 the contents of which are provided by the remote and local identification fields 72 and 74 in data base 70 (FIG. 3B). These fields 132 and 134 identify the virtual
10 circuit through which the message is being transmitted. Since each receiving unit may have several virtual circuits, even between itself and the same node or device server, the fields 132 and 134 are used to identify the specific virtual circuit for which the message received
15 over communications link 16 is intended. If the message is a START virtual circuit message from device server 24, with reference to FIG. 6B, the destination virtual circuit identification field 132 contains a "0", and the source virtual circuit identification field contains an
20 identification that is assigned by the device server 24.

The virtual circuit header 130 also includes a message type field, an M flag and an R flag the contents of which are provided by fields 76, 78, and 80 in the virtual circuit data base 70 (FIG. 3B). The header 30
25 also includes message sequence and acknowledgement sequence numbers taken from counters 82 and 84, and a field 136 which identifies the number of session slots that may be included in a data field 140. The contents of session number field 136 is used only in connection
30 with a RUN virtual circuit message (FIG. 6C) described below. In a START virtual circuit message, the data field 140 contains information used by the recipient in setting up the virtual circuit, and in a STOP virtual circuit message (FIG. 6D) the data field contains

information as to the reason the virtual circuit is being stopped.

A virtual circuit message (FIG. 6A) ends in an error check field 142 which contains a cyclic redundancy check word used to verify that the message was received without errors.

With reference again to FIG. 4, after the device server 24 transmits the START virtual circuit message, it shifts from a HALTED state to a STARTING state.

10 Similarly, when the node 34 receives the START virtual circuit message, it shifts from a HALTED state to a STARTING state and responds with either a START virtual circuit message, indicating that it will support and participate in the virtual circuit, or a STOP virtual
15 circuit message, indicating that it will not support the virtual circuit. In either case, the node 34 retrieves the contents of the source virtual circuit identification field 134 from the message and stores it in the remote identification field 72 in its virtual circuit data
20 base 70 as the device server's identification of the virtual circuit.

With reference again to FIG. 6B, if the node responds with a START virtual circuit message, the node uses the contents of the source virtual circuit
25 identification field 134 from the device server's START virtual circuit message as the contents of the destination virtual circuit field 132 in its responding START virtual circuit message. The node also provides the contents of the source virtual circuit identification
30 field 134 (FIG. 6A) as its identification code for the virtual circuit. The device server retrieves the contents of this field, stores it in the remote identification field in its data base 70 for this virtual circuit and thereafter uses it in the destination virtual

circuit identification field 132 of future messages transmitted over the virtual circuit.

With reference to FIG. 6D, if the node responds with a STOP virtual circuit message, it also provides a source
5 virtual circuit identification code and the data field 140 identifies the reason that the node will not support the virtual circuit; one such reason may be that the node 34 is currently supporting other virtual
10 circuits and has insufficient resources to provide support for another virtual circuit. If the node 34 transmits a STOP virtual circuit message to the device server, both the node and the device server return to the halted state. The device server may then attempt to
15 establish a virtual circuit to another node connected to a service provider that provides the desired service or inform the user that the service is not available if no other node provides the service.

In addition, the user may determine, after the device server has transmitted a START virtual circuit
20 message, that it does not need to use the particular service. This is indicated in FIG. 4 by a USER HALT directed at the STARTING state of device server 24. If that occurs, the device server transmits a STOP virtual circuit message (FIGS. 6A and 6D) to node 34. Both
25 virtual circuit state machines 60 and 64 then return to the halted condition.

If node 34 responds with a START virtual circuit message, and if no USER HALT occurs at the device server 24 in the starting state, both the state
30 machines 60 and 64 (FIG. 3A) shift to a RUNNING state. In this condition, the device server 24 and node 34 can transmit RUN virtual circuit messages depicted in FIGS. 6A and 6C. In this condition, the data field 140 contains session slots which are described below (FIGS. 5

and 7A through 7D). The number of session messages is identified in the session number field 136 in the virtual circuit header 130 (FIG. 6A). In the session messages, the device server 24 and node 34 transmit service
5 information between the service users 12 and service providers 14 (FIG. 1), more specifically the service information is transmitted between service state machines 66 and 68.

When the user no longer needs a service, it
10 disconnects from the service, and, if no other users are using the virtual circuit, a USER HALT condition exists. If the state machines 60 and 64 are both in the RUNNING state, the device server may transmit a STOP virtual circuit message (FIGS. 6A and 6D) to node 34 and return
15 to the HALTED state.

As has been noted above, when the virtual circuit state machines 60 and 64 are in the running state, the device server 24 and node 34 can transmit RUN virtual circuit messages which include session slots. Using the
20 session slots, the session state machines 66 and 68 are established, and service data and status information are transmitted between the service user and provider. When the service user no longer needs the service, the session state machines may then be abolished, thereby terminating
25 the service session. The session slots which are depicted in FIGS. 7A through 7D. By means of the session slots, the session state machines 66 and 68 shift among various states as depicted in FIG. 5.

With reference to FIG. 5, the session state
30 machine 66 of device server 24 has five states, including a HALTED state, a STARTING state, an ABORT START state, a RUNNING state, and a STOPPING state. The session state machine 68 of node 34 has four states, including a HALTED state, a STARTING state, a RUNNING state, and a STOPPING

state. Initially, both state machines 66 and 68 are in a halted state, and when a user requests a particular service, the device server transmits a START session message, in a virtual circuit message through the virtual
5 circuit.

With reference to FIG. 7A, the format of a session message includes a session header 150 which includes destination session identification field 152, a source session identification field 154, a byte count field 156,
10 a session slot type field 158, a credits field 160. A session data field 162 carries information for establishing and abolishing a session and service session data and status information. The destination and source session identification fields 152 and 154 are used in the
15 identical manner as the destination and source virtual circuit identification fields 122 and 124 (FIG. 6A) described above. The contents of these fields are stored in and taken from the remote and local identification fields 102 and 104 in the session data base 100 (FIG. 3C).

20 The contents of the byte count field 156 identify the length of the session data field 162, and are taken from the byte count field 110 in the data base 100. The contents of the session slot type field are taken from field 112 in the session data base, and identify the type
25 of message being transmitted. As noted above, five types of session slots may be transmitted. The contents of the CREDITS field 160 is taken from the local credits field 114 in session data base 100 (FIG. 3C) to identify the number of slots that are available in the data
30 buffer 106 for any response. When a unit receives a session slot, the contents of the CREDITS field may be stored in the remote credits field 116 of the session database 100 to indicate the amount of space in the data

buffer 106 available for a subsequent session slot transfer.

In the START session slot (FIG. 7B) the session data field 162 provides information required by the session state machine at node 34 for setting up the session. Such information may include, for example, the type of service required, thereby identifying the service provider which is to engage in the service session, and size of the data buffer 106 set aside for the session in the device server.

After the device server 24 transmits the START session slot, the session state machine 66 shifts to a STARTING state. After the node receives the START message, the node's session state machine 68 shifts to the STARTING state and the node responds with either a START session slot, after which the state machine 68 shifts to the RUNNING state, or a REJECT session slot (FIG. 7D), after which the state machine shifts to the HALTED state. In either case, the node 34 supplies its session identification code in the source session identification field 154 (FIG. 7A). If the node responds with a REJECT session slot, the credits field 160 also contains the reason that it is rejecting the service session. Such reasons may include, for example, that the node is unable to provide the service because of inadequate resources, or that the node or the service provider is shutting down.

With reference again to FIG. 5, the device server 24 may receive the START session or REJECT session slot from node 34 when it is in either the STARTING state or in an ABORT START state. The device server's session state machine 66 shifts from the STARTING state to the ABORT START stage if a user disconnect request, indicating that the user does not want to use the service which it

previously selected, is received from a user before either a START session or a REJECT session slot is received from node 34 . If the session state machine 66 is in the ABORT START state, and if a REJECT session slot is received, the session state machine 66 merely returns to the HALTED state. However, if a START session slot is received, the device server 24 transmits a STOP session slot over the virtual circuit to the node causing its session state machine 68 to return to the HALTED state.

10 In either case, both session state machines 66 and 68 return to the HALTED state.

However, if the device server's session state machine 66 is in the STARTING state, and if the device server receives a START message from the node 34, its session state machine shifts to the RUNNING state. In that state, and with the node's session state machine 68 in the RUNNING state, RUN session slots including service data and status information may be transmitted back and forth between device server 24 and node 34 over the virtual circuit. With reference to FIGS. 7A and 7C, the session data field 162 in such messages contains user and service provider data and status information.

15
20

After the operator determines that it no longer requires access to the service, the user requests disconnection from the service and both state machines 66 and 68 shift to the STOPPING state and then return to the halted state.

25

FIG. 8. depicts the timings of various messages transmitted between a device server 24 and node 34, in response to the service circuit timer 92 and the R flag 80, the R flag in the server being set or cleared by the R field in a virtual circuit header in a message from the node. FIG. 8 also illustrates the timings of the retransmissions of various messages that are not received

30

by the server and node, in response to the re-transmit
timers 90 in each unit. Specifically, when a node
receives a message from the device server, it will
respond with a message having the R field being either
5 set or clear depending on whether or not the message
includes data. If the message does not include any data,
the R field is clear, and the node may thereafter
transmit a second message to the device server which
includes data. The second message has a set R field.
10 This is illustrated in time (E). As shown at times (F)
and (G), the second message may cross with another
message from the device server, which may transmit a
message when its data waiting flag is set, when its
service circuit timer times out. The R field being set
15 forces the data waiting flag to be set, thereby enabling
the server to send an acknowledging message to the node,
whether or not it actually has data to send. The R field
being clear enables the node to thereafter send another
message with data, thereby removing the constraint on the
20 node that it send messages, and therefore data, only when
it receives messages from the server.

A detailed description of the operation of
network 10 (FIG. 1) will now be presented. Periodically,
each node 34 connected into network 10 transmits a
25 multicast service advertising message (FIG. 2A)
identifying the particular services which are available
through it. All of the device servers receive the
advertising message and establish a service directory
(FIG. 2B) identifying the services which are available
30 and the nodes and ratings of the services available
through each node. The available services may be
displayed to the operators of the service users from the
device servers' directories.

When a service user 12 requires the use of a service, its device server 24 determines which node 34 provides the service and which has the highest service rating for that service. The device server 24 then
5 determines whether it has a virtual circuit 58 between itself and that node. If no such virtual circuit exists, the device server transmits a START virtual circuit message (FIGS. 6A and 6B) to the node to attempt to establish a virtual circuit over communications link 16.
10 If the node responds with a START virtual circuit message, the virtual circuit has been established and the device server 24 then establishes a session between the virtual circuit 58 and the user 12 requesting the service. It will be appreciated that, if a virtual
15 circuit already exists to the required node, the device server need not set up another virtual circuit, but instead may proceed to the next steps and use the existing virtual circuit for communications. With some communications links, such as a link conforming to the
20 Ethernet protocol, it may be desirable to limit the number of users using a single virtual circuit, since the length of each virtual circuit message is limited. Accordingly, even if a virtual circuit is already established between the device server and node, it may be
25 desirable to establish additional virtual circuits if too many users are already using one virtual circuit.

After the virtual circuit is established, the device server 24 then may transmit virtual circuit messages through the virtual circuit including session messages
30 multiplexed from several service sessions. A session begins with a START session slot (FIGS. 7A and 7B) transmitted to device server 24 to node 34 identifying the required service, in an attempt to establish a service session over the virtual circuit 58 with the

service provider 14 which provides the required service. If the service session is established, service data and status information can be transmitted in RUN session messages.

- 5 The rate at which the device server can transmit virtual circuit messages over the communications link 16 is limited by the server circuit timer 92 (FIG. 3B) thereby allowing messages for other virtual circuits to be multiplexed over the communications link.
- 10 Furthermore, the data waiting flags 86 in the respective device server and node inhibit them from transmitting virtual circuit messages through the virtual circuit until it has information transmit. Thus, neither the device server nor the node transmits any messages over
- 15 the virtual circuit 58 unless information is available to be transmitted, and then no more often than permitted by the server circuit timer.

When a user finally determines that it no longer requires a service, the session can be abolished by

20 device server transmitting a STOP session slot. If all of the sessions are abolished for a virtual circuit, the device server 24 may then abolish the virtual circuit by the transmission of a STOP virtual circuit message.

If additional service users 12 require services

25 provided by a node 34, their session slots may be carried by the virtual circuit messages transmitted over virtual circuit 58. Thus, a single virtual circuit message transmitted over the virtual circuit 58 can contain messages between a number of users 12 and service

30 providers 14. By establishing a virtual circuit between the device server 24 and node 34, and multiplexing session slots onto a single virtual circuit message, the number of acknowledgement messages may be reduced to one response message for each virtual circuit message. This

reduction results in, not only a reduced amount of message traffic over the communications link since only the virtual circuit messages between nodes and device servers are acknowledged, and not the messages between particular service users and service providers. In the past, each message between service users and service providers had been acknowledged by a separate acknowledgement message, which not only increased message traffic over the communications link, but also required processing activities by the service user and service provider that are not required in the network constructed in accordance with the invention.

The foregoing description has been limited to a specific embodiment of the invention. It will be apparent, however, that the invention can be practiced in networks having diverse basic construction than is disclosed in this specification, with the attainment of some or all of the advantage of the invention. Therefore, it is the object of the appended claims to cover all such variations and modifications as come within the true spirit and scope of the invention.

CLAIMS :

- 1 1. A local area network for interconnecting service users
and service providers, including a plurality of device
server units each of which connects to a service user and
a plurality of nodes each connected to a service provider
5 and a communications link to effect communications
between the nodes and device server units to enable said
service users and service providers to communicate,
 - A. each node including means for transmitting a
service advertising message over said communications link
10 to identify the services provided by the service
providers connected thereto, and
 - B. each device server unit including:
 - i. advertising message receiving means connected to
said communications link for receiving said service
15 advertising message;
 - ii. service table means for establishing a service
table including a plurality of entries each including a
node field and a service field identifying a node and a
service provided by the service providers connected to
20 the node; and
 - iii. selection means responsive to an operator
requesting a service for using said service table to
select the node and service provider to provide the
25 requested service.
- 30 2. A node for connection in a local area network, the
network interconnecting service users and service
providers and including a plurality of device server
units each of which connects to a service user and a
plurality of nodes each connected to a service provider
and a communications link to effect communications
between the nodes and device server units to enable said
service users and service providers to communicate the
node including means for transmitting a service

1 advertising message over said communications link to
identify the services provided by the service providers
connected thereto.

5 3. A device server unit for connection in a local area
network, the network interconnecting service users and
service providers and including a plurality of device
server units each of which connects to a service user and
a plurality of nodes each connected to a service provider
and a communications link to effect communications
10 between the nodes and device server units to enable said
service users and service providers to communicate, each
node including means for transmitting a service
advertising message over said communications link to
identify the services provided by the service providers
15 connected thereto, the device server unit including:

 A. advertising message receiving means connected to
said communications link for receiving said service
advertising message;

20 B. service table means for establishing a service
table including a plurality of entries each including a
node field and a service field identifying a node and a
service provided by the service providers connected to
the node; and

25 C. selection means responsive to an operator
requesting a service for using said service table to
select the node and service provider to provide the
requested service.

30 4. A local area network for interconnecting service users
and service providers, including a plurality of device
server units each of which connects to a service user and
a plurality of nodes each connected to a service provider
and a communications link to effect transfers of messages

1 between the nodes and device server units to enable said service users and service providers to communicate, each said device server unit including:

A. timing means for generating a timing signal at the end of a predetermined timing interval;

5 B. data storage means for accumulating data from said service users for transmission to said service providers;

10 C. data waiting flag means connected to said data storage means and having a set condition when data is stored in said data storage means and a reset condition; and

15 D. message transmission means connected to said communications link, said timing means, said data storage means and said data waiting flag means for transferring data from said data storage means over said communications link in response to the data waiting flag means having a set condition if said timing means has generated the timing signal.

20 5. A local area network as defined in claim 4 wherein said node includes node message transmission means connected to said service providers and said communications link for transmitting messages to a device server unit in response to the receipt of a message from a device server unit.

25 6. A local area network as defined in claim 5 wherein each message originating at a node includes a response requested flag field having selected values, said device server unit further having message receiving means for receiving messages, said message receiving means
30 including means for setting said data waiting flag means in response to the response requested flag field of a message having a predetermined value.

1 7. A local area network as defined in claim 6 wherein a
message from said node to said device server unit
includes a data field for transferring data from said
node to said device server unit, said node transmitting a
5 message having a response requested flag field having
said predetermined value when said data field includes
data and having the other predetermined value when said
data field is empty, said node message transmission means
thereafter being inhibited from sending another message
10 when said response requested flag has said first
predetermined value until the node receives a subsequent
message from said communications link, said node message
transmission means being enabled to send another message
at any time when said response requested flag has said
15 other predetermined value.

8. A device server unit for connection in a local area
network for interconnecting service users and service
providers, the network including a plurality of device
20 server units each of which connects to a service user and
a plurality of nodes each connected to a service provider
and a communications link to effect transfers of messages
between the nodes and device server units to enable said
service users and service providers to communicate, each
25 said device server unit including:

A. timing means for generating a timing signal at
the end of a predetermined timing interval;

B. data storage means for accumulating data from
said service users for transmission to said service
30 providers;

C. data waiting flag means connected to said data
storage means and having a set condition when data is
stored in said data storage means and a reset condition;
and


1 D. message transmission means connected to said
communications link, said timing means, said data storage
means and said data waiting flag means for transferring
data from said data storage means over said
communications link in response to the data waiting flag
5 means having a set condition if said timing means has
generated the timing signal.

9. A device server unit as defined in claim 8 wherein a
message from a node to said device server unit includes a
10 response requested flag field having selected values,
said device server unit further including message
receiving means for receiving messages from said
communications link, said message receiving means
including means for setting said data waiting flag means
15 in response to a response requested flag field having
predetermined ones of said selected values.

10. A node for connection in a local area network for
interconnecting service users and service providers, the
20 network including a plurality of device server units each
of which connects to a service user and a plurality of
nodes each connected to a service provider and a
communications link to effect transfers of messages
between the nodes and device server units to enable said
25 service users and service providers to communicate, each
said device server unit including means for transmitting
messages in response to the setting of a data waiting
flag means normally indicating that it has data to send
on the timing out of a timing means, said node including
30 message receiving means for receiving messages from the
communications link and message transmission means for
transmitting a message including a data field and a
response requested flag field having selected values, the

1 values of said response requested flag field being
established in response to whether the data field is
empty.

5 11. A node as defined in claim 10 further comprising
inhibiting means connected to said message transmission
means and said message receiving means for inhibiting
said message transmission means from transmitting a
message until said message receiving unit has received a
message, and enabling means responsive to the value of
10 the last transmitted response requested flag field for
disabling said inhibiting means when the value of said
response requested flag field indicated that the data
field of the message was empty.



LOCAL AREA NETWORK 10

FIG. 1

SERVICE ADVERTISING MESSAGE

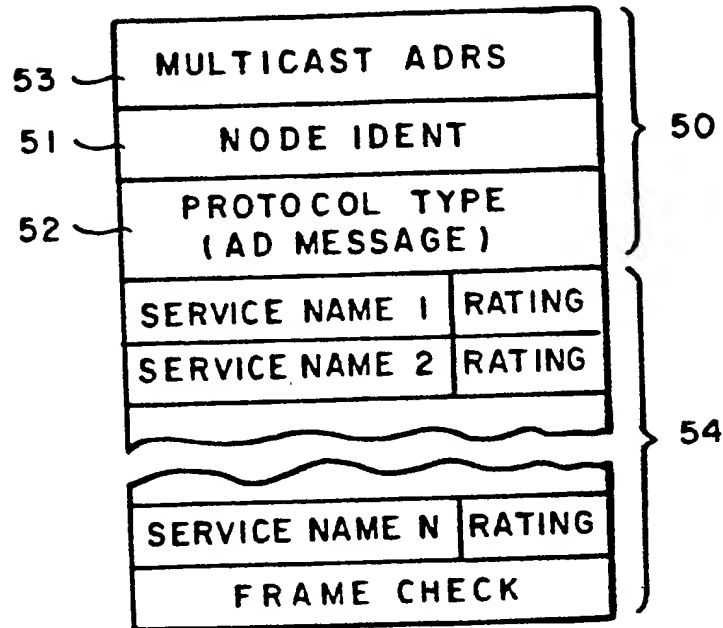


FIG.2A

SERVICE DIRECTORY

NODE	SERVICES	RATING
A	DISK A	
B	DISK A	
C	PHONE LINES	
D	PROC A PROC B TAPE	
E	PROC C PROC D TAPE	

FIG.2B

VIRTUAL CIRCUIT 58




FIG. 3A

3.11
SERVICE PROVIDERS 14

0163577

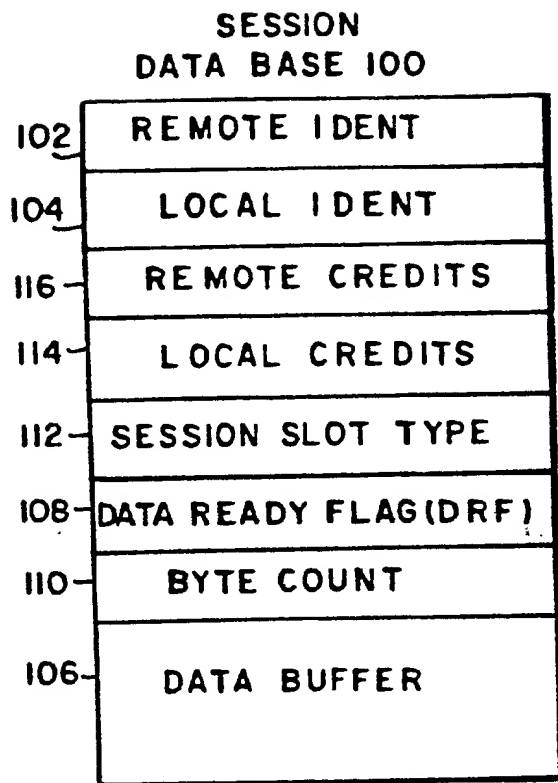


FIG.3C


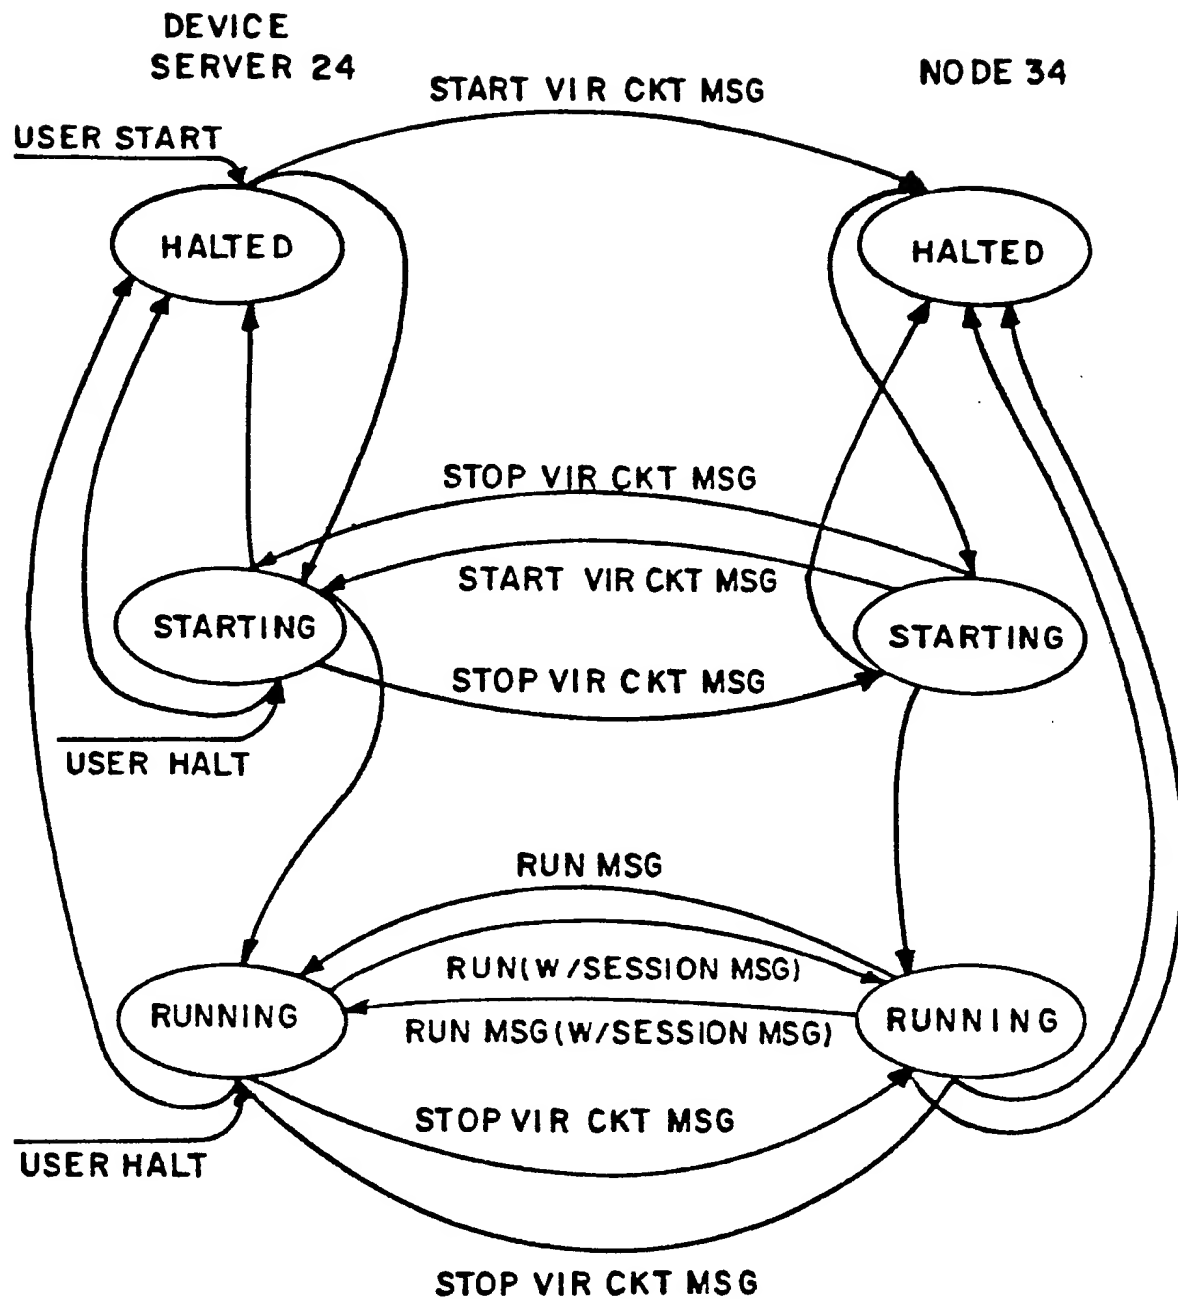
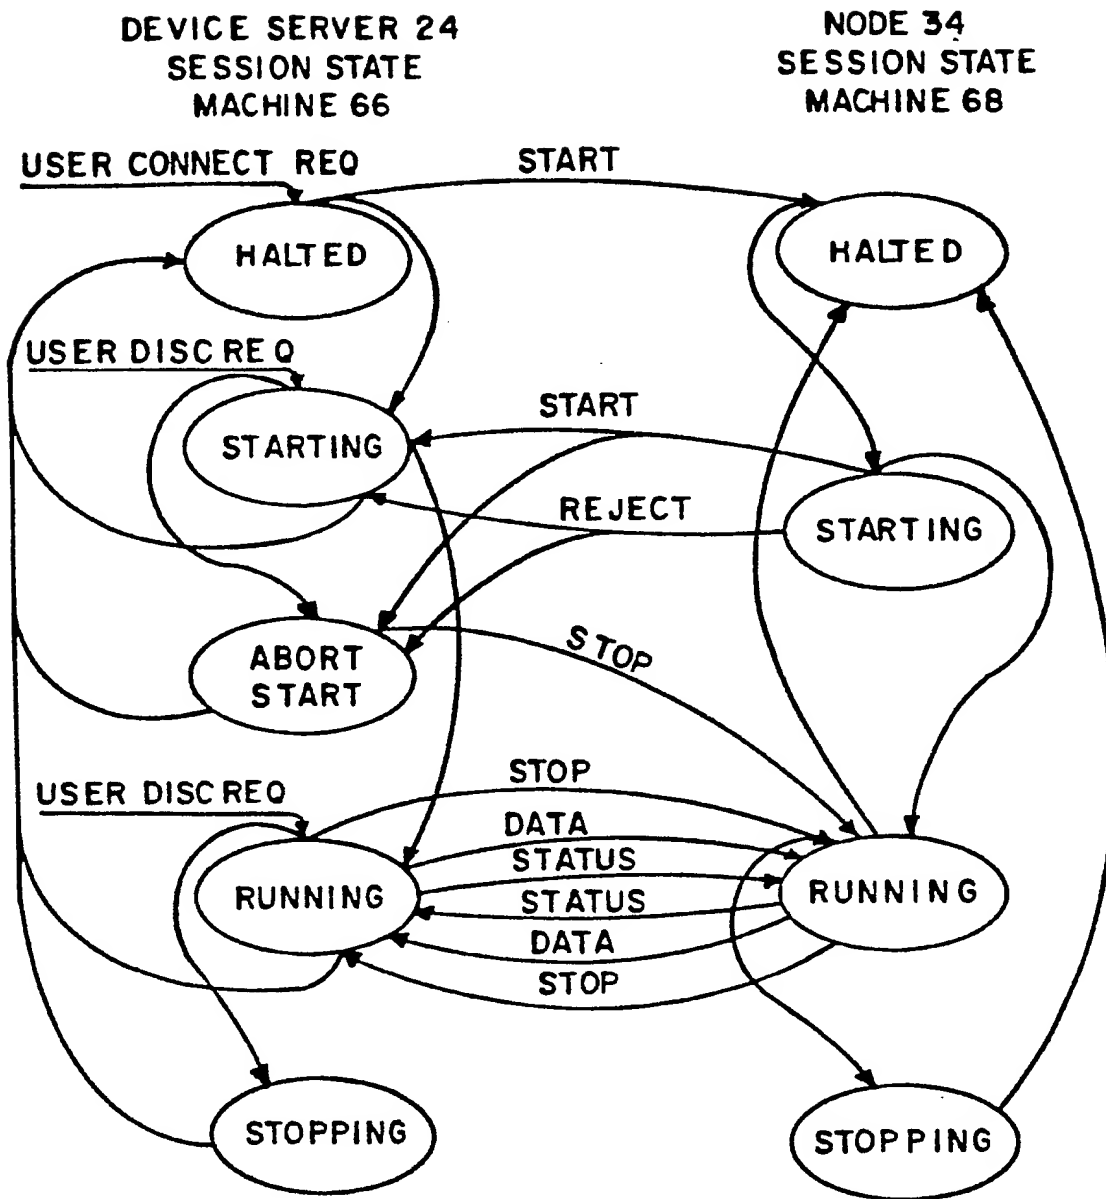


FIG.3B



VIRTUAL CIRCUIT
STATE DIAGRAM

FIG. 4



SESSION STATE
DIAGRAM

FIG.5



8.11

RUN VIRTUAL CIRCUIT MESSAGE

MSG TYPE = Ø

FROM DEVICE SERVER { DEST VIR CKT ID ← NODE ASSIGNED
SRC VIR CKT ID ← DEV SERV ASSIGNED

FROM NODE { DEST VIR CKT ID ← DEV SERV ASSIGNED
SRC VIR CKT ID ← NODE ASSIGNED

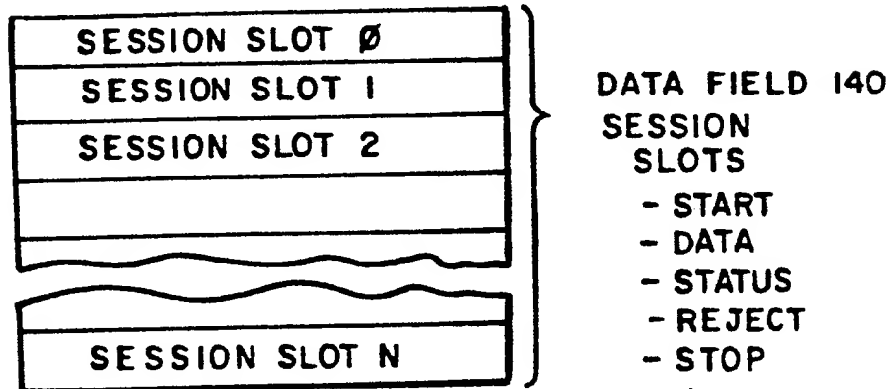


FIG.6C

STOP VIRTUAL CIRCUIT MESSAGE

MSG TYPE = 2

FROM DEVICE SERVER { DEST VIR CKT ID ← NODE ASSIGNED
SRC VIR CKT ID ← DEV SERV ASSIGNED

FROM NODE { DEST VIR CKT ID ← DEV SERV ASSIGNED
SRC VIR CKT ID ← NODE ASSIGNED




FIG.6D

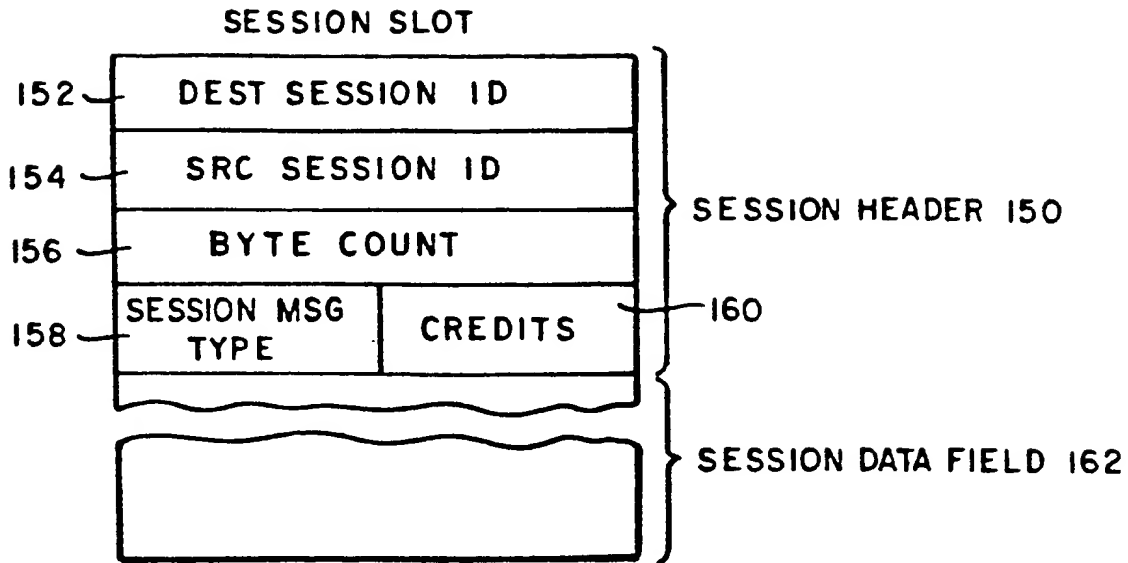


FIG.7A

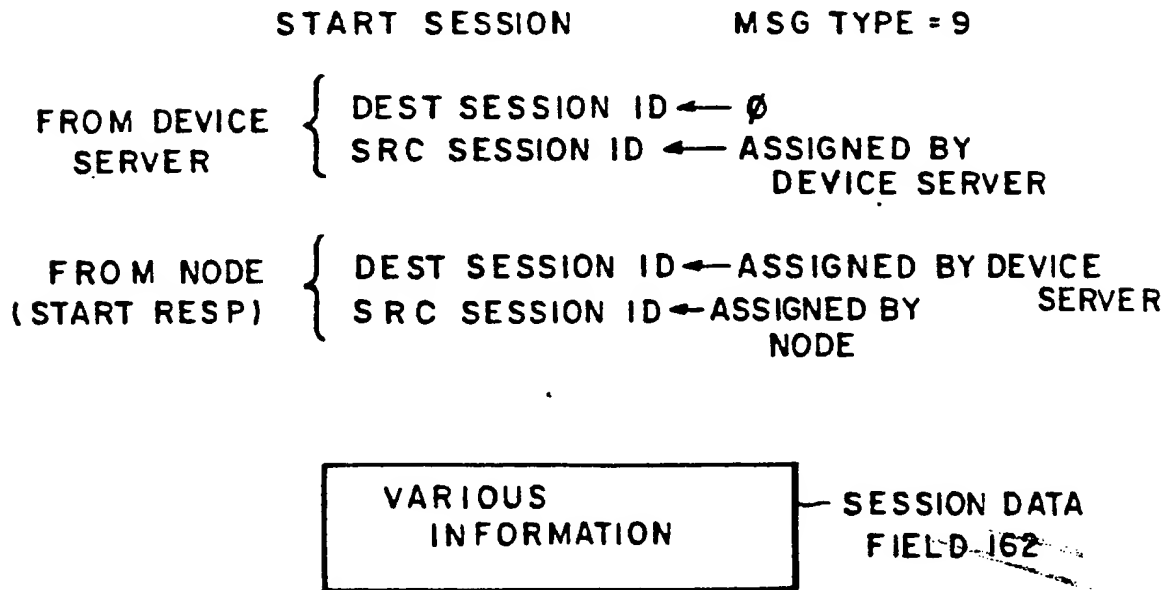


FIG.7B

10.11

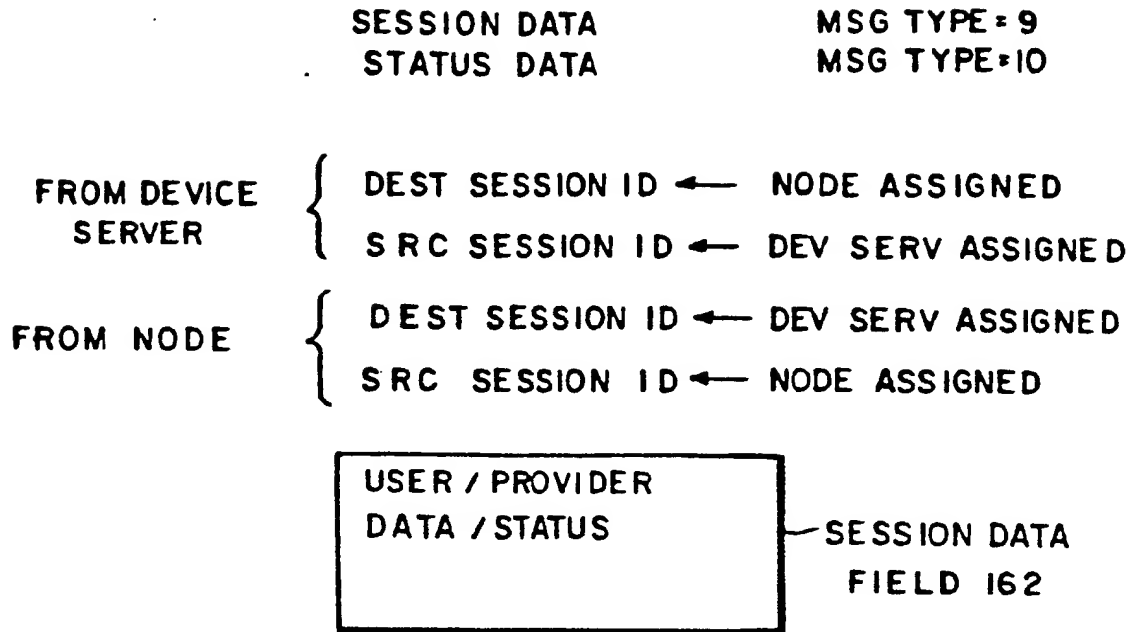


FIG. 7C

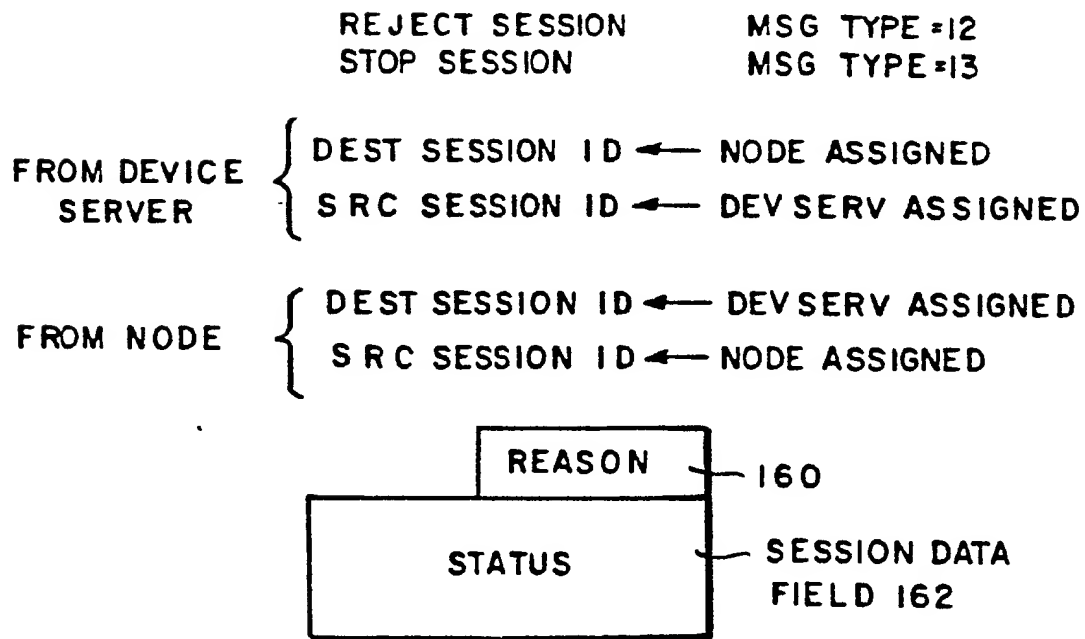


FIG. 7D




FIG. 8